RESOURCE RESERVATION in ADVANCE in HETEROGENEOUS NETWORKS with PARTIAL ATM INFRASTRUCTURES

Alexander Schill, Sabine Kühn, Frank Breiter
Dresden University of Technology

Abstract

Resource reservation in advance (ReRA) enables scheduling and allocation of resources at an early stage in time. This way, the availability of resources can actually be guaranteed for the point in time when the resources are needed. As opposed to that, current reservation protocols such as RSVP perform an "immediate" reservation without advance scheduling. They can therefore suffer shortage of resources with subsequent rejection of applications.

This paper describes our new approach for providing resource reservation in advance in ATM networks. As a foundation, we first present experiences of IPng and RSVP over ATM. Based on this work, we then discuss major design issues of an appropriate ReRA solution. Important aspects are the signalling model for ReRA, the admission control strategy, the duration of connections, and various other time-related parameters. We also discuss our ongoing implementation of ReRA as an extension of RSVP and outline directions for future research in this area.

1: Introduction

ATM networks have recently gathered great interest in research and practice. According to the current UNI 3.1 (User Network Interface) specification and the UNI 4.0 specification of the ATM forum [3,4], important quality of service (QoS) characteristics are supported. However, there are two major problems with this approach:

(1) Heterogeneous networks: Most current and future networks are not purely based on ATM but consist of a mix of various technologies (ATM, FDDI, Ethernet etc. with switches, hubs, routers etc. for network interconnection). Therefore, higher-level protocols on top of ATM such as IP are necessary to enable heterogeneous interoperability. Moreover, reservation protocols at this level such as RSVP (resource reservation protocol; e.g. for reserving resources within routers) are also necessary to achieve QoS guarantees.

(2) Drawbacks of immediate reservation: Current resource reservation facilities, both within ATM and on top of it, perform a so-called "immediate" reservation: They reserve resources just at the point in time when they are actually needed. If resources cannot be granted at this time, QoS cannot be guaranteed as desired or important applications such as broadband videoconferences even have to be rejected. An advance scheduling of reservations could help to solve this problem.

The major contributions of this paper to address these problems are as follows:

(1) Experiences with IP/IPng and RSVP over ATM: In order to address the first problem, we present an implementation of the IP/IPng (IP next generation) and RSVP (resource reservation) protocols at OSI layer 3 on top of ATM. This enables extended addressing capabilities in Internet environments, enhanced QoS specifications, and associated resource reservation in heterogeneous networks.

(2) Resource reservation in advance (ReRA): In order to address the second problem, we present a new approach for implementing advance scheduling of reservations. This way, the availability of resources can actually be guaranteed for the point in time when the resources are needed. Important design decisions discussed are the signalling model for ReRA, the admission control strategy, the duration of connections, and various other time-related parameters. We also discuss our ongoing implementation of ReRA as an extension of RSVP.

The paper is organized as follows. In section 2, we discuss the current situation in terms of standards and common protocols in the context of ATM. In particular, IP/IPng over ATM is discussed. Moreover, related work concerning resource reservation in advance is outlined. Section 3 presents our implementation and experiences with RSVP over ATM. Section 4, the main part of the paper, discusses concepts and design decisions of our resource reservation in advance approach. In section 5, implementation aspects of ReRA as an RSVP extension are presented as a synergy of the former considerations.

2: Background and Related Work

As of today, ATM networks mainly offer adaptation layer AAL5 supporting available bit rate (ABR) traffic, and, in some cases, AAL1 offering a constant bit rate (CBR). Although applications could in principle directly

0-8186-7780-5/97 $10.00 © 1997 IEEE

use the AAL interface, the dominance of heterogeneous networks and the required migration of existing applications onto ATM makes higher-level protocols necessary.

A typical approach standardized is IP over ATM [14]: Applications just use conventional TCP/IP or UDP/IP, and IP communication is mapped onto virtual ATM connections. IP addresses are mapped onto ATM addresses by an ATMARP server (address resolution protocol), and IP packets are segmented into ATM cells and reassembled at the receiver's site transparently. The ATMARP approach requires that communication between different IP subnetworks is performed via routers, even if two communicating nodes are attached to the same physical ATM network. The emerging NHRP (next hop resolution protocol) will solve this problem by enabling direct ATM shortcuts whenever possible. The use of a MARS (multicast address resolution server) [20] also enables the implementation of IP multicast over ATM. As opposed to the current overlay model (overlay of IP routing on top of an ATM network), the peer model [1] is a direct algorithmic mapping of all network layer addresses onto ATM addresses, and there is only one kind of signalling, i.e. ATM signalling as defined by the PNNI specification (private network node interface) [19]. This leads to an environment with combined ATM switches and routers. An alternative solution for mapping existing applications onto ATM is LANE (LAN Emulation) [12, 13], which only enables the interconnection of homogeneous LANs via ATM. Therefore, we do not consider LANE in this paper. Currently, MPOA (multiprotocol over ATM) [1] is also discussed as a kind of integration approach; it supports LAN protocols, IP, IPX and other protocols over ATM.

With these solutions, however, ATM is merely a "large pipe" with significant throughput, but without any quality of service guarantees. Only best-effort traffic based on ABR can be supported as the given protocols do not provide QoS specification and resource reservation facilities. Therefore, problems with realtime applications such as videoconferences with shared applications, remote video monitoring, time-constrained bulk transactions etc. arise. An early, sender-oriented reservation protocol addressing these issues was ST-II [8]. More recently, RSVP (resource reservation protocol) [5] has observed growing interest. Basically, RSVP enables the reservation of bandwidth and other resources at the time when a resource-intensive communication scenario is initiated. This way, actual QoS guarantees can be given. RSVP coexists with IP or IPng (also known as IPv6 [11], as a successor of IPv4); it operates only during call setup and maps reservation requirements directly onto underlying networks (such as ATM) and network components (such as routers). Due to its importance in the upcoming Integrated Services Internet [7], RSVP has been selected as a basis of our implementation work, too.

Even with RSVP, however, reservations can of course fail if resources are not sufficiently available at the time when they are needed. Therefore, some first approaches towards providing resource reservation in advance facilities have been developed.

In summary, it can be stated that most of the known papers briefly describe more general problems to be solved: [10] concentrates on resource partitioning; [15] describes a possible exchange of advance reservation information using the ST-II [8] reservation protocol. Dealing with the Integrated Services Internet, [9] presents an explicit admission control algorithm for predictive service. Moreover, [17] already introduces a possible ReRA architecture and [16] compares the suitability of the Internet-based reservation protocols ST-II and RSVP to support ReRA. However, implementations have obviously not reached a mature stage yet.

Our approach differs from these solutions with respect to two major aspects: First, we consider quite a large number of design parameters in terms of signalling, admission control, and timing aspects and aim at achieving a large degree of flexibility. Secondly, our solution is directly based on RSVP and ATM as emerging standards.

3: RSVP over ATM

As originally conceived, the Internet, and also solutions based on IP over ATM offer only a very simple kind of QoS: best effort. Hence, before realtime applications like video conferencing can be used, the Internet has to be modified to support real time service. This is only possible by resource reservation. A protocol like RSVP is crucial for the negotiation of the corresponding reservation parameters and for guaranteeing quality of service (QoS) characteristics based on explicitly reserved network, memory and processing resources [6]. It is of particular importance in heterogeneous networks with partial ATM infrastructures. With the assistance of RSVP, the specification of QoS (Quality of Service) and traffic parameters for dedicated flows can be performed by the applications on top of the IP/IPng protocol stack. The usage of the new Internet protocol in conjunction with RSVP allows a more efficient handling and classification of packets as belonging to an already reserved flow.

3.1: RSVP - An IP Signalling Protocol

With the usage of RSVP [18], recently proposed as a Request for Comments, the Integrated Services Internet [7] especially supports realtime applications with guaranteed, predictable and controlled end-to-end performance across networks. The main task performed by RSVP is signalling of resource requirements at connection setup time. To be precise, RSVP does not actually reserve or allocate resources but rather indicates reservation requests to the

underlying systems.

Moreover, the protocol offers a flexible handling of heterogeneous receivers as well as an adaptation to dynamically changing multicast groups. The transmission of RSVP control information is implemented by encapsulating RSVP packets into IP or IPng packets. The basis of a reservation is a detailed description of flow traffic and QoS characteristics. In accordance with this fact, RSVP defines the so-called flow (describing the traffic and the required QoS parameters) and filter specification (identifying the flow for which reservations have been performed). All reservation and state information are managed in soft states in all nodes which facilitates the renegotiation of resource reservations within RSVP.

Even though the reservation is receiver-oriented, the initiator of a reservation is the sender, which informs appropriate receivers about the characteristics of the flow to be sent. Figure 1 shows the main protocol elements.

In contrast to the RSVP model, the QoS setup time in

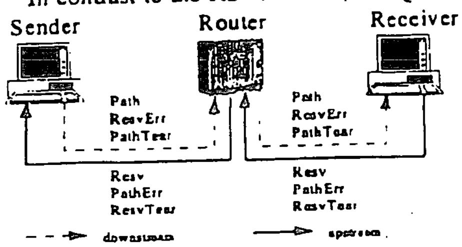

Figure 1: Exchange of RSVP messages

ATM is directly combined with the connection setup. Considering both the ATM and the RSVP service classes (of IETF), the latter are characterized by delay, whereas, in ATM, classes of bit rates (ABR, CBR etc.) are specified. Moreover, ATM allows both variable bit rate (burst, peak) and constant bit rate for transmission of non bursty traffic as opposed to RSVP. The traffic parameters of ATM accordingly comprise Maximum Burst Size, Sustainable Cell Rate and Peak Cell Rate. The mapping of the particular RSVP traffic parameters is discussed in conjunction with an assignment of the service classes in [2].

3.2: RSVP-IP/IPng-ATM Protocol Stack

Within our development work, an RSVP implementation on top of ATM has been completed under the Digital Unix operating system. As part of the IP/IPng convergence modules (CM), and hence part of the ATM subsystem, the RSVP specific API supports real time handling and transport of IP/IPng flows (figure 2).

The API is able to receive reservation information from a local RSVP daemon. Based on information contained in a given flow specification, a new ATM VC reservation is performed after completing the mapping of service classes and parameters. The classification of the IP/IPng flows belonging to a dedicated reserved reservation (virtual channel) is done in the IP/IPng CM, too.

As a basis for the implementation of ReRA, we use the full RSVP/IP(IPng)/ATM protocol stack which we have completed in cooperation with Digital Equipment.

The following section describes the general approach

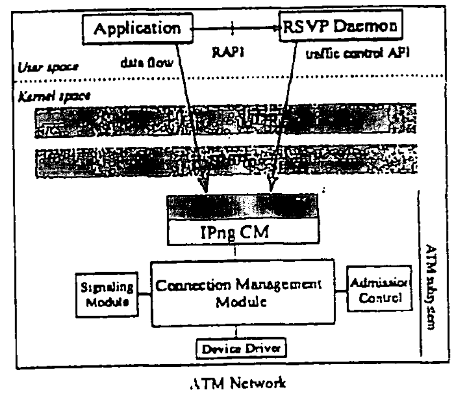

Figure 2: ATM subsystem and RSVP

to realize an extension of the "immediate" reservation mechanisms of RSVP and ATM with advance scheduling and planning techniques.

4: Resource Reservation in Advance

4.1: An Introduction

Considering, for instance, the „real world", activities such as a meeting in a conferencing room are scheduled for a specific time. With this example, reservation means that limited resources are reserved a certain time in advance to get an assurance that the resources are available at the requested time. The definition of reservation in current QoS based networks like ATM covers another kind of semantics: reservations are being performed in conjunction with a connection establishment; this means that reservations are done at the time when the network resources are actually needed. This is an element of uncertainty, because in the worst case the required resources are used by other applications and therefore the application request must be rejected. However, using ReRA, a future lack of resources during actual reservations can be avoided.

From our point of view, it shall be possible to schedule applications, e.g. an important videoconference with several partners, for a given time in the future. The system shall then calculate and virtually reserve the required resources for that time, however without immediately blocking them. In order to grant the required QoS, it is necessary to provide mechanisms of resource reservation in advance.

The advantage of introducing a ReRA concept is also reasonable in QoS based high performance networks, in which a lot of users have to share the same resources.

4.2: ReRA Connection Parameters

Comparing the negotiation of immediate and advance reservation, one of the differences is the specification of additional parameters for setting up a reservation in advance. As the resources for advance reservations will be reserved for a given time in the future, one of the necessary parameters is the time at which the resources will be needed.

Furthermore, it seems to be meaningful to specify a duration for ReRA-connections which describes how long the reservation will be alive. Based on this information the admission control is able to recognize whether the reservation is overlapping with others or not. On one hand, specifying a reservation duration results in a more effective utilization of bandwidth (discussed in more detail in 4.4.2). On the other hand, many applications are not always able to specify a concrete duration for their reservations.

4.3: Signalling Model Description

The differences between the ImRe and ReRA communication model is shown in figure 3. For describing the extended communication to set up a resource reservation in advance, we have introduced time marks at which messages are sent or actions are performed.

To specify how much of the resource capacities have to be reserved, the client issues a REQUEST at the time $t_{request}$. It also specifies the points in time that define beginning (t_{start}) and duration (d_{res}) of the reservation. After successfully completed negotiation and admission control, the service provider CONFIRMs the reservation (at the time $t_{confirm}$). Moreover, each reservation is assigned an identifier for later access. The service provider RELEASEs the request in the case of rejection.

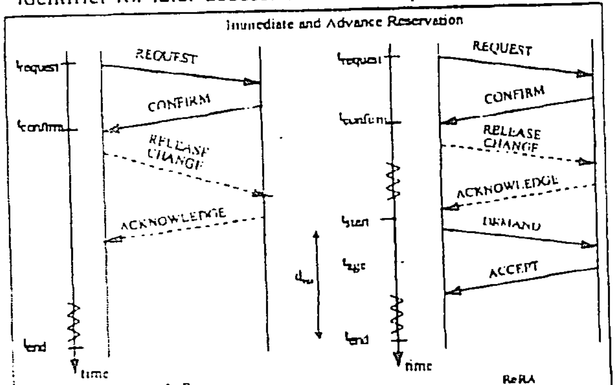

Figure 3: Communication flow with ImRe and ReRA

After a successful negotiation, the users might want to re-negotiate the communication resources. This covers the cancellation (RELEASE) or an increase or decrease (CHANGE) of the requested resources. If the changes are acceptable or the reservation was successfully released the service provider ACKNOWLEDGEs it. Re-negotiation concerns not only the traffic and QoS parameters of the reserved connection but also the starting point and duration. A re-negotiation can be performed until t_{end} is reached. To realize a more flexible handling of changes, policy data within a CHANGE request have been introduced to indicate that the reservation should either be deleted or kept in the intermediate nodes if the modification fails. Another possible policy is that any rejection carries suggestions for alternative reservations: (1) If a reservation cannot be granted for the whole duration requested, the admission control has to determine the possible (shorter) duration. (2) It is also possible that the system suggests a new start time at which the requested resources will be available or (3) it just provides the currently available resources (i.e. less than requested) for the given time.

The decision, which of the policies should be used for ReRA or ImRe is a responsibility of the network administrator or provider. In our solution such information is exchanged with the user whenever a request fails.

As the interval between $t_{confirm}$ and t_{start} could cover a long time, it is possible that intermediate nodes may be unplugged from the network for a short time, may crash or the topology may change. The occurrence of such kinds of failures is not different from failures within traditional QoS-based systems during the data transmission phase [17]. However, the handling of these failures differs dependent on the applied reservation system. The handling of failures before the data transmission starts via a ReRA connection requires special attention. In chapter 5 an example of such a special handling is given for RSVP.

In particular, the transition to the data transmission phase can be considered as a special case. When the start time is reached, the DEMAND for using the reservation has to be performed within t_{start} - t_{age}. If the sender is not sending a DEMAND by then, the reservation will be aged and all state information corresponding to it is removed. In contrast to immediate reservation, ReRA therefore requires additional communication at t_{start}. We prefer this solution because last changes should be possible and resources are only reserved if they are really needed. If an application exceeds its requested duration without a recently requested prolongation, the system handles it as an immediate reservation.

4.4: Call Admission Control

4.4.1: Desired Features

The design of a suitable ATM traffic control is considered as a fundamental challenge for the success of the ReRA mechanism. The primary role of traffic control is to protect the network and the user in order to achieve predefined network performance objectives. Dealing with

advance reservation, an additional role of traffic control is to optimize the use of network resources for the purpose of achieving: (1) sufficient network efficiency, that means the effectiveness in utilizing bandwidth, (2) a minimal blocking probability and (3) a fair handling of the advance and the immediate reservations. Obviously, it seems to be difficult to combine these features in such a way that the Call Admission Control (CAC) will achieve all these goals in an optimal way. Therefore, we are investigating influence parameters (e.g. bandwidth awarding, connection parameters) and their effects on the characteristics of the CAC.

4.4.2: Impacts on the Admission Control Features

These and other influences concern the awarding of the resources between the two types of reservations. That means precisely the kind of policing of resource awarding and management that should be used between ImRe and ReRA. Another interesting point which shall be investigated is the sending of immediate versus deferred acknowledgements after receiving a setup request. Corresponding to this fact, the introduction of a granularity of time intervals is investigated, influencing the determination of starting time of a future reservation and of its duration.

Moreover, the knowledge of the connection duration has an important impact on the effectiveness of CAC and also on the acceptance of applying reservations. Finally, a minimum and maximum advance notice time for submitting reservations should be specified.

4.4.3: Awarding of Bandwidth

Basically, there are the following policies to handle the two different kinds of reservations: sharing and partitioning. Sharing the resources allows the awarding of the whole bandwidth e.g. in the case of an aggressive access to ReRA connections. A partitioning of the bandwidth for instance for ImRe and ReRA is heading towards an equitable handling of the single reservations as the bandwidth is dedicated to each reservation type. But, assuming an adverse allocation of resources between ImRe and ReRA, it can result in an inefficient utilization of bandwidth caused by the two boundaries.

4.4.4: Immediate and Deferred Acknowledgements

To optimize the bandwidth utilization and to reduce blocking probability, an investigation of immediate versus deferred acknowledgement of reservations is interesting. Generally considered, there are two models based on the time when scheduling decisions are made. In the case of immediate acknowledgement, scheduling decisions are made as soon as a request arrives. The other model is based on delayed acknowledgements, only applicable to reservations in advance. Every request has an associated decision point which is given by the system. Focusing on a reduced blocking probability, for example, the following simple heuristics can be used. We assume, that the bandwidth requests of all reservations with a synchronous starting point are sorted in an ascending way. Accepting all reservations beginning with the smallest resource request until the whole bandwidth is awarded, results in an reduced blocking probability. Obviously, assuming synchronous starting points of an amount of reservations is an abstract model and only conceivable with relatively coarse-grained time intervals. In the future, we also have to investigate and develop reasonable heuristics to achieve a better utilization of resources by a special handling of requests according to their attributes (e.g. bandwidth requirements).

4.4.5: Duration of Connections

The specification of the connection duration will influence the effectiveness in utilizing bandwidth dependent on the selected resource awarding strategy. If the bandwidth between ImRe and ReRA is shared, then the bandwidth utilization is optimized by specifying the duration for all ImRe and ReRA connections.

While this allows the system to make more efficient use of resources by explicitly limiting the duration of connections, it may be unacceptable to some applications. Using partitioning between ImRe and ReRA (figure 4), a definition of the duration is necessary only for ReRA connections to make decisions about future resource allocation. To be precise, in the case of static partitioning, the immediate partition treats any new requests exactly as usual. The advance partition must instead use a different mechanism to test new requests for admission.

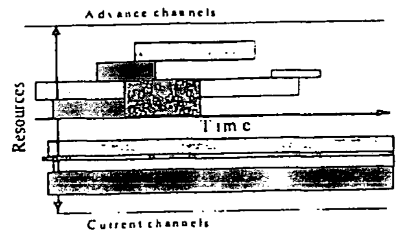

Figure 4: Resource partitioning

4.4.6: Granularity of Starting Points and Duration

Basically, an introduction of a granularity of time intervals has an effect on the administration of the reservations on the various intermediate nodes. Using a granularity results in a more effective mutual coincidence of reservation start and end points, that means connections could only be allowed to be pre-reserved at certain times. From this point of view, a rough granularity is more convenient. Considering typical applications, reservations are generally made in steps of five or up to 15 minutes [15].

For example (figure 5), a mapping of the event C onto the next starting time of the next granularity unit results in

an inaccuracy of up to 1 time slot G; with a reasonable choice of C, this should usually be acceptable for most applications.

In spite of these facts, the approach has the side-effect that a clock synchronisation throughout the network must be assumed, as different times on different nodes may result in askew starting times due to different roundings.

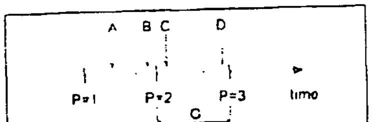

Figure 5: Usage of time granularity

4.4.7: Minimum and Maximum Advance Notice

The maximum notice as well as the minimum advance notice should be reasonable values with which a reservation can be submitted (e.g. for less than three hours and not more than 3 months) if such limits were imposed by the provider. As the value of the maximum advance notice has an impact on the management overhead of the admission control, it should not exceed a reasonable value. Moreover the minimum advance notice is also used to differ between ImRe and ReRA requests.

4.5: An Approach of ReRA in ATM

For implementing ReRA in an ATM environment, the following aspects have to be considered: (1) how to realize the necessary communication for advance reservation in ATM; (2) which components have to be changed and (3) how to realize the call admission control.

Although, there is more than one communication mechanism in ATM possible for exchanging ReRA information, the extension of the signalling protocol seems to be the most general approach. Considering ReRA in a private ATM network, an extension and modification of the signalling concerns both the User-Network-Interface and the Network-Network-Interface signalling protocol.

Recognizing ReRA in ATM means not only the transmission of ReRA information (4.5.1) and the decision on the acceptance or rejection of a new call (4.5.2). Because ReRA differs from immediate reservations, a possible QoS (load sensitive) routing and reactions to switch failures should also be considered. In the case of ImRe, the state and load information are exchanged between the switches by conveying appropriate cells as a part of the routing protocol of the PNNI [19]. A ReRA-QoS routing algorithm would allow the network capacity to be engineered efficiently, by virtue of its ability to take into account the future network state in making routing decisions. This may enable more effective operation and can reduce blocking probability.

4.5.1: ReRA Signalling

Introducing an extended signalling to negotiate ReRA parameters is realized in modifying existing and defining new message types. Modifying messages comprises the addition of newly defined and extended information elements (IE). These IE are necessary to convey the following parameters: starting time, duration and policy information for the requested ReRA. As the IEs will be handled as optional, the message types can be further used for immediate reservations in their original meaning.

The negotiation of immediate and advanced reservation are nearly the same, so that the **SETUP**[1] message is always used to *REQUEST* the needed resources. Concerning ImRe, sending a **SETUP** message results in a connection establishment. As opposed to that, it requests the setup of an advance reservation with ReRA. Based on the information of the included ReRA IE, the system then calculates and virtually reserves the required resources for that time, however without immediately blocking them. Furthermore, the **CONNECT, RELEASE_COMPLETE, CALL_PROCEEDING** and the **RELEASE** message are used without any modifications. Except the former message, all other messages can be used within the time $t_{request} - t_{end}$.

After a successful negotiation of a reservation, it is possible to modify the starting time and duration of the ReRA connection, whereas the renegotiation of traffic and QoS parameters is allowed in the intermediate phase, too. Analogous to the recent ATM signalling, it is not possible to change any other parameter than traffic and QoS parameters in the usage phase without releasing the old and establishing the new connection. New message types were necessary due to an additional communication making modifications and demands of the already requested resources (ReRA) possible:

Hence, the *CHANGE* message is comparable to the **SETUP** message; in relation to the conveyed parameters, however, it has a different meaning. To be more precise, a *CHANGE* message is used to modify an existing request whereas a **SETUP** message should be used to initiate a new request. In the *CHANGE* message policy data can be used to indicate that the request should be either deleted or kept in the intermediate nodes if the modification fails. Concerning the specification of the policy data for *CHANGE*, it is sufficient to provide one IE in which each policy is represented by setting the corresponding bit.

The transition from the intermediate phase to the usage phase will be done by sending a *DEMAND* message. In analogy to the *CHANGE* message, the *DEMAND* message looks like the **SETUP** message.

---

[1] Words in bold face indicate ATM signalling messages and italic ones indicate ReRA messages.

Most of the IEs in the *DEMAND* message are optional as well as in the *CHANGE* message. Hence, it is possible to create a „very short" *DEMAND* message, which results in a decreased processing time in the end systems and the switch. We call this behaviour a "fast setup". Moreover, by setting up the appropriate IEs, it is allowed to make last minute changes by the user, sending a *DEMAND*. The messages which are sent to indicate an *ACCEPTion* are similar to the acknowledgement of the **SETUP** message.

4.5.2: CAC Decisions

As we consider ReRA only as an additional service, the impacts or influences on ImRe shall be small. So we decided that the underlying policy of admission control should be awarding bandwidth fairly between both kinds of reservations based on an optimized bandwidth utilization: a partitioning with a moveable boundary [10]. This means, each reservation has its own partition, but with the option of using available resources of the other partition if the own partition is completely allocated. To avoid a sharing across the whole bandwidth, we have introduced watermarks to protect e.g. the advance partition. Moreover, the immediate reservation can be handled as a special case of ReRA with starting time 0. In the case of a fully utilized ImRe bandwidth, the ReRA part can be used by reservations with start time "now" only with the maximum duration of minimum advanced notice. A part of the immediate reservation partition should be protected by a watermark, too. Considering the immediate partition, the following characteristics are notable: (1) it can be used for infinite immediate reservations, and (2) only for advance reservations if the advance partition is saturated.

The advance partition can be used: (1) for advance reservations and (2) limited immediate reservations.

Figure 6 shows an example that immediate reservations could be limited by the system. Obviously,

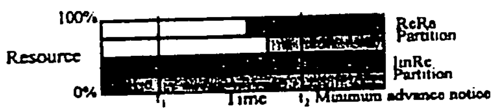

Figure 6: Filling gaps caused by ReRA with limited ImRe

this was also introduced to realize a more efficient utilization by filling gaps caused by the reservations in advance. This refers to bandwidth which is available in the time interval t_1 to t_2 and cannot be used by advance reservations in the advance partition. It can therefore be allocated to ImRe with limited duration.

If an infinite ImRe was requested and only resources within a limited duration are available, the user could be informed by the system as part of a rejection of the request according to the policies described in chapter 4.3.

5: ReRA in RSVP

In this chapter essential changes and extensions of RSVP to support ReRA will be described. In particular, this includes the negotiation of ReRA information, and the behaviour of RSVP in case of a temporary router failures before the starting time of the advance reservation is reached.

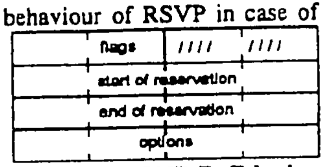

Figure 7: RSVP ReRA class

The ***negotiation*** of ReRA information is done by sending Path and Resv messages carrying an additional instance: a new RSVP class which is shown in figure 7.

As the ReRA object is an additional and optional object in an RSVP message, a result of this solution is that the semantics of the messages will be preserved for ImRe. This implies, that no new messages have to be defined in RSVP; instead the new RSVP class has been introduced which includes the points in time when the reservation starts and finishes. The flags can be used, for example, to distinguish between a normal Resv message and a Resv message (the latter complies with a DEMAND message, chapter 4.3), which is send at the starting point to indicate that the user wishes to use the reserved resources. Moreover, to change the end of an already requested reservation, in the ReRA object carried by a Resv message a special flag is set and the options field carries the new end time.

Moreover, to realize a ***complete extension of the RSVP specification*** to support ReRA we have performed the following steps:

- the extension of the messages (described above) and the corresponding extension of the message processing rules,
- the extension of the state control blocks (Path-, Reservation- and Traffic Control blocks) to support the pseudo hard state concept (see below) and to hold the new reservation parameters,
- the extension of the specific RSVP interfaces (e.g. RAPI, traffic control API, chapter 4.3) and the introduction of new functions, for example, as part of the RAPI, to allow: (1) a LEAVE of a session without releasing the soft states. This is used to inform the RSVP daemon that the application is not actual available for the whole time until the starting point is reached (for instance 3 months). The application should save session information to be able to re-register by calling the session function. (2) As we allow more than one reservation in one session (which can be distinguished by the start and end time of the reservation) the TEARDOWN function can be called to remove a specific reservation.

To avoid that soft states which hold the ReRA information will be deleted caused by temporary router failure or a longer termed unplugged endsystem we use a modified refreshing mechanism and the concept of pseudo hard states and proxy routers. Considering the soft state

approach of RSVP (chapter 3), the periodic sending of Path and Resv messages (default: every 30 seconds) produces a considerable overhead. The refreshing of the states is controlled by the refresh time R and the lifetime L. The relationship between R and L is expressed by the following equation: $L \geq (K + 0.5) \cdot 1.5 \cdot R$ (K-1 successive messages may be lost without state information being deleted). From our point of view, L should increased in such a way that reservations are insensitive to temporary router failures. This can be achieved by an adequate increase of R at the time the ReRA connection was requested and a reduction of R step by step until the starting time is reached. At this time, R is set to the default value. It should be emphasized that concrete values of R and K have to be determined based on future measurements in our experimental environment.

As the advance notice of an advance reservation request can take a long time (e.g. 3 months), the actual availability of endsystems have to be considered. If an endsystem is not actually available (for instance a PC which was turned off), this results in missed Path and Reserve refresh messages and in a loss of the corresponding soft states of the neighboured router. In order to avoid this, the routers which are situated immediately behind the sender and the receiver, respectively, should react if no new refresh message arrives. Because these routers take over the function of the endsystems in sending Path and Resv refresh messages, they will be called proxy routers.

Moreover, each router has to determine if the next host is the receiver downstream or the sender upstream, itself. Is one of cases fulfilled, this results in holding the soft states the router. More precisely, these soft states are kept and only removed if the user sends a teardown message or if the starting time of the reservation is already exceeded. Therefore, the states are called pseudo hard states and so we are able to deal with unplugged endsystems.

6: Conclusions

This paper has presented the current situation in terms of standards and common protocols in the context of ATM. In particular, IP and IPng over ATM were discussed. Due to its importance in the upcoming Integrated Services Internet, RSVP has been selected as the basis of our implementation work. As a main contribution of this paper, a general model of advance reservation was described and an enhancement of RSVP and ATM was presented to support reservation in advance. Special emphasis was put on the problem of signalling the appropriate ReRA information within the integrated services Internet and the ATM network. For mapping the general model onto RSVP and ATM, differences in handling failure situations must be considered; they closely depend on the current handling of routing changes and neighbour losses concerning traditional reservations.

The presented concepts for resource reservation in advance are also under refinement and we expect to have further implementation experiences soon.

7: References

[1] Alles, A., „ATM Internetworking". 1995

[2] Schill, A., Kühn, S., Breiter, F., Internetworking over ATM: Experiences with IP/IPng and RSVP, 7th Joint European Networking Conference, Budapest, Hungary, 1996

[3] ATM-Forum, „UNI Signaling 4.0", ATM Forum 94-1019R6, Work in Progress, 1995

[4] ATM-Forum, „Traffic Managment Specification Version 4.0", ATM Forum 95-0013R6, Work in Progress, 1995

[5] Braden, R., Zhang, L., „Resource ReSerVation Protocol (RSVP) -- Version 1 Functional Specification", Work in Progress, 1996

[6] Campbell, A., Coulson, G., Hutchison, D., „A Quality of Service Architecture", Computer Communication Review, Vol. 1, No. 2, April 1994, pp. 6-27

[7] Integrated Services Model, RCF 1633, 1995

[8] Delgrossi, L., Herrtwich, R.G., Hoffmann, F., „An Implementation of ST-II for the Heidelberg Transport System", Internet-working: Research and Experience, Vol. 5, 1994, pp. 43-69

[9] Degermark, M., Köhler, T., Pink, S., Schelén, O., „Advance Reservation for Predictive Service", 5th International Workshop on Network and Operating System Support for Digital Audio and Video, Durham, April 1995

[10] Ferrari, D., Gupta A., Ventre, G., „Distributed advance Reservation of real-time connections", 5th International Workshop on NOSSDAV, Durham, April 1995

[11] Hinden, R., Deering, S., „IP Version 6 Addressing Architecture", Work in Progress, 1995

[12] Jeffries, R., „ATM LAN Emulation", Data Communications, No. 9, 1994, pp. 95-100

[13] ATM-Forum, „LAN Emulation over ATM Version 1", 1995

[14] Laubach, M., „Classical IP and ARP over ATM", RFC 1577, January 1994

[15] Reinhardt, W., „Advance Reservation of Network Resources for Multimedia Applications", Proceedings of the 2nd IWACA94, Heidelberg, 1994

[16] Reinhardt, W., „Advance Resource Reservation and its Impact on Reservation Protocols", Technical Report, 1995

[17] Wolf, L.C., Delgrossi, L., Steinmetz, R., Schaller, S., Wittig, H., „Issues of Reserving Resources in Advance", IBM European Network Center Heidelberg, TR 43.9503, 1995

[18] Zhang, L., Deering, S., Estrin, D., Shenker, S., Zappala, D., „RSVP: A New Resource Reservation Protocol", IEEE Network, Sept. 1993, pp. 8-18

[19] ATM-Forum, „PNNI Spezification Version 1.0", 1996

[20] Armitage, G., „Support for Multicast over UNI 3.0/3.1 based ATM Networks", RFC 2022, 1996